arXiv:1604.06515v1 [stat.ME] 22 Apr 2016

A weighted edge-count two-sample test for multivariate and object data

Hao Chen, Xu Chen, and Yi Su*

Abstract

Two-sample tests for multivariate data and non-Euclidean data are widely used in many fields. Parametric tests are mostly restrained to certain types of data that meets the assumptions of the parametric models. In this paper, we study a nonparametric testing procedure that utilizes graphs representing the similarity among observations. It can be applied to any data types as long as an informative similarity measure on the sample space can be defined. The classic test based on a similarity graph has a problem when the two sample sizes are different. We solve the problem by applying appropriate weights to different components of the classic test statistic. The new test exhibits substantial power gains in simulation studies. Its asymptotic permutation null distribution is derived and shown to work well under finite samples, facilitating its application to large datasets. The new test is illustrated through an analysis on a real dataset of network data.

Keywords: nonparametric test, unequal sample sizes, permutation null distribution, similarity graph.

*Hao Chen is an Assistant Professor at the Department of Statistics at University of California, Davis. Xu Chen is a Master's student at the Department of Statistics at Duke University. Yi Su is a PhD student at the Department of Statistics at University of California, Davis.

1 Introduction

Two-sample testing is a fundamental problem in statistics. Due to the increasing richness of data in both dimension and complexity, this problem is encountering new challenges. Nowadays, many applications involve the test on data in high dimensions [de la Sierra et al., 2011, Feigenson et al., 2014] or even on non-Euclidean data, such as image data and network data [Eagle et al., 2009, Kossinets and Watts, 2006]. Parametric approaches can be applied to multivariate data under certain assumptions while their power decreases quickly as the dimension grows unless strong assumptions are made to facilitate the estimation of the large number of (nuisance) parameters, such as the covariance matrix. In this work, we study a nonparametric testing procedure that works for both multivariate data and object data.

Nonparametric testing for two sample differences has a long history and rich literature; see Gibbons and Chakraborti [2011] for a survey. Friedman and Rafsky [1979] proposed the first practical test that can be applied to data with arbitrary dimension. They used pairwise distances among the pooled observations from both samples to construct a minimum spanning tree (MST), which is a spanning tree that connects all observations with the sum of distances of edges in the tree minimized. The test statistic is the number of edges that connect nodes (observations) from different samples. We call this test the *edge-count test* for easy reference.

The edge-count test is not restricted to the MST. It can be applied to any similarity graph where more similar observations are more likely to be connected. Friedman and Rafsky [1979] also considered k-MSTs. A k-MST is the union of the 1st, ..., and kth MSTs, where the ith MST is a spanning tree connecting all observaitons that minimizes the sum of distances across edges subject to the constraint that this spanning tree does not contain any edge in the 1st, ..., (i-1)th MST(s). They showed that the edge-count test on a 3-MST is usually more powerful than that on a 1-MST. Schilling [1986] and Henze [1988] used k-nearest neighbor graphs where each observation is connected to its k closest neighbors. More recently, Rosenbaum [2005] proposed to use the minimum distance non-bipartite pairing (MDP). This divides the N observations into $N/2$ (assuming N is even) non-overlapping pairs in such a way as to minimize the sum of $N/2$ distances within pairs.

For an odd N, Rosenbaum suggested creating a pseudo data point that has distance 0 to all observations, and later discarding the pair containing this pseudo point. This way of constructing the graph can be extended to k-MDPs as well, where a k-MDP is defined similarly to a k-MST. Besides these common ways to construct the similarity graph, the graph can also be provided by domain experts directly [Chen and Zhang, 2013].

The rationale of the edge-count test is that, if the two samples are from different distributions, observations would be preferentially closer to those from the same sample than those from the other sample. Thus edges in the similarity graph would be more likely to connect observations from the same sample. The test rejects the null hypothesis of equal distribution if the number of between-sample edges is significantly *less* than what is expected under null. Maa et al. [1996] showed that the edge-count test based on MST constructed on Euclidean distance is consistent against all alternatives for multivariate data.

However, in practice, sample sizes range from tens to thousands, or somewhat larger. We call these sample sizes *practical sample sizes*. Chen and Friedman [2016] found that, when the dimension of the data is moderate to high, for practical sample sizes, the edge-count test is effective for locational alternatives but can have low power for scale alternatives when the Euclidean distance is used to construct the similarity graph. The authors proposed a new test statistic on the similarity graph that works for both locational and scale alternatives under practical sample sizes. We call this test the *generalized edge-count test* for easy reference.

Chen and Friedman [2016] recommended to use the generalized edge-count test when there is no clue on the type of alternatives. However, when the alternative is locational, the edge-count test is recommended as it in general has higher power than the generalized edge-count test under such circumstances.

In this work, we addressed another problem for the edge-count test when the sample sizes of the two samples are different. Taking the edge-count test on the 5-MST constructed on Euclidean distance for testing the mean difference of two 20-dimensional Gaussian distributions as an example, we found that, starting from the equal sample size scenario, the power of the edge-count test *decreases* when one sample size is doubled and the other

sample size keeps the same (see in Section 2 for more details). This is counter-intuitive as increasing the sample size adds more information and we would expect the power of the test to increase. This weird phenomenon indicates that the edge-count test statistic is not well defined.

To address this problem, we propose a modified version of the edge-count test. The idea is that, when the sample sizes are different, it is more difficult to form an edge within the sample with a smaller sample size than that for the sample with a larger sample size. So we give within-sample edges different weights according to which sample they are from instead of treating them equally in the edge-count test. The new test works properly under unequal sample sizes and we call it the *weighted edge-count test*.

When the sample sizes are different, under locational alternatives, the weighted edge-count test is more powerful than the edge-count test and the generalized edge-count test. Hence, the weighted edge-count test and the generalized edge-count test can be used complementarily with one mainly for locational alternatives and the other for more general alternatives.

The rest of the paper is organized as follows. The problem of the edge-count test under unequal sample sizes is explored in Section 2. The weighted edge-count test is proposed and studied in Section 3. Its power is examined in Section 4 and its asymptotics are studied in Section 5. We illustrate the weighted edge-count test through an analysis on a real dataset of network data in Section 6. In Section 7, we discuss in more details the relation between the weighted edge-count test and the generalized edge-count test.

2 The problem

In this section, we explore what happens to the edge-count test under unequal sample sizes. As an illustration example, we randomly draw m observations from distribution $F_1 = \mathcal{N}(\mathbf{0}, I_d)$ and n observations from distribution $F_2 = \mathcal{N}(\boldsymbol{\mu}, I_d)$, $\|\boldsymbol{\mu}\|_2 = 1.3$, $d = 50$, and call them Sample 1 and Sample 2, respectively. Here, $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ denotes a multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix Σ, and I_d is a $d \times d$ identity matrix. We use the common notation $\|\cdot\|_2$ to denote L_2 norm. The two distributions are chosen such that the test has moderate power. We consider the following two scenarios:

- Scenario 1: $m = n = 50$.
- Scenario 2: $m = 50,\ n = 100$.

So for scenario 2, one sample size keeps the same and the other sample size is doubled. We applied the edge-count test on k-MST constructed on the Euclidean distance to the simulated data. Intuitively, the test should have a higher power in scenario 2 than in scenario 1 because there are more observations in scenario 2. We estimated the power of the test by the fraction of trials that the test rejected the null hypothesis at 0.05 significance level in 1,000 trials. The results are shown in Figure 1.

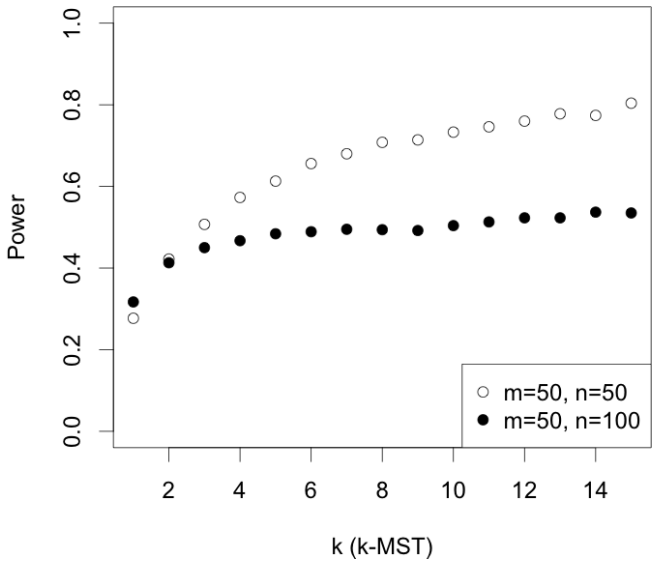

Figure 1: The fraction of the trials (out of 1,000) that the edge-count test on k-MST rejected the null hypothesis of equal distribution.

The choice of an optimal k is an open question. We here show the results from 1-MST to 15-MST. To our surprise, we see from Figure 1 that the edge-count test has lower power in scenario 2 under all k-MSTs except for 1-MST. Why does this happen?

Before exploring in details, we first introduce some notations. Let $N = m + n$ be the total sample size. We pool observations from both samples and index them by $1, \ldots, N$. Let G be an undirected similarity graph on the pooled observations (nodes) with no multi-edge, i.e., there is at most one edge between any two nodes. The graph can be a k-MST, a

k-MDP, etc. We use G to refer to both the graph and its set of edges when the vertex set is implicitly obvious. The symbol $|\cdot|$ is used to denote the size of the set, so $|G|$ is the number of edges in G. Let R be the number of edges in G that connect observations between the two samples, i.e., the number of between-sample edges. we work under the permutation null distribution, which places $1/\binom{N}{m}$ probability on each of the $\binom{N}{m}$ permutations of the sample labels. When there is no further specification, we denote by $\mathbf{P}$, $\mathbf{E}$, $\mathbf{Var}$, $\mathbf{Cov}$ probability, expectation, variance, and covariance, respectively, under the permutation null distribution.

We next explore in details on what happens in the edge-count test. We focus on the test based on 5-MST (similar patterns are observed for other k-MSTs where $k > 1$). First of all, we check whether adding more observations does make R further smaller than its null expectation $\mathbf{E}(R)$. In each run, we calculate $\mathbf{E}(R) - R$ for scenario 1 and denote it by D_1. We then add 50 more observations randomly drawn from F_2, re-construct the 5-MST based on the 150 observations, calculate $\mathbf{E}(R) - R$ based on the new 5-MST, and denote it by D_2. We check whether D_2 is larger than D_1 in general. Figure 2 shows the boxplots of D_1 and D_2 separately, as well as their differences $(D_2 - D_1)$, from 1,000 simulation runs.

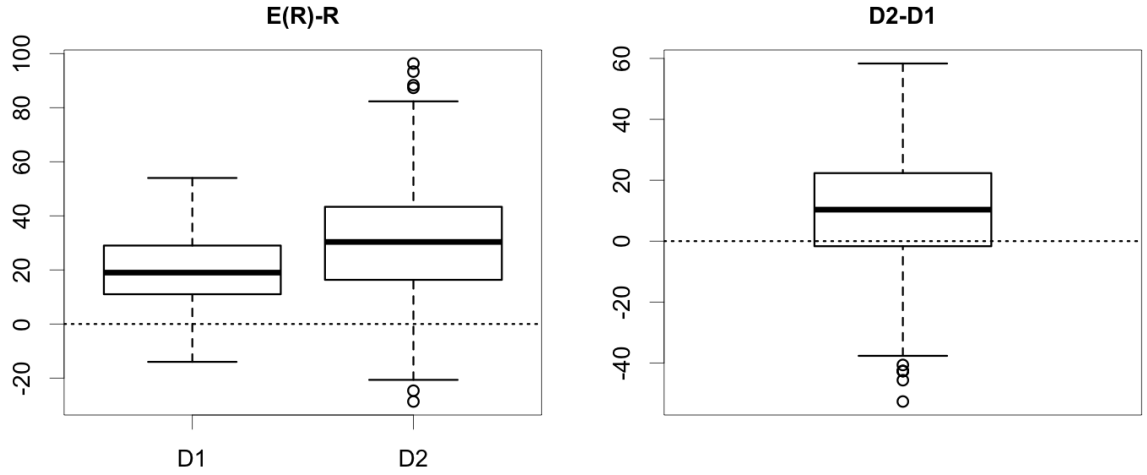

Figure 2: Boxplots of $\mathbf{E}(R) - R$ before and after adding 50 more observations, as well as their difference (after − before), from 1,000 simulation runs. The horizontal dashed line is at level 0.

We see that $\mathbf{E}(R) - R$ is in general positive under both scenarios. This is expected since the two distributions differ in the mean and the observations tend not to find observations from the other sample to be similar. The boxplot of the difference between D_2 and D_1 shows that $\mathbf{E}(R) - R$ indeed becomes larger in general when more observations are included. This also complies with what we would expect.

To quantify how further R is from its null expectation $\mathbf{E}(R)$, we need to compare $\mathbf{E}(R) - R$ to its standard deviation under null. Figure 3 shows the boxplots of the standard deviations of R before (sd_1) and after (sd_2) adding the 50 more observations, as well as their ratio (sd_2/sd_1), from 1,000 simulation runs.

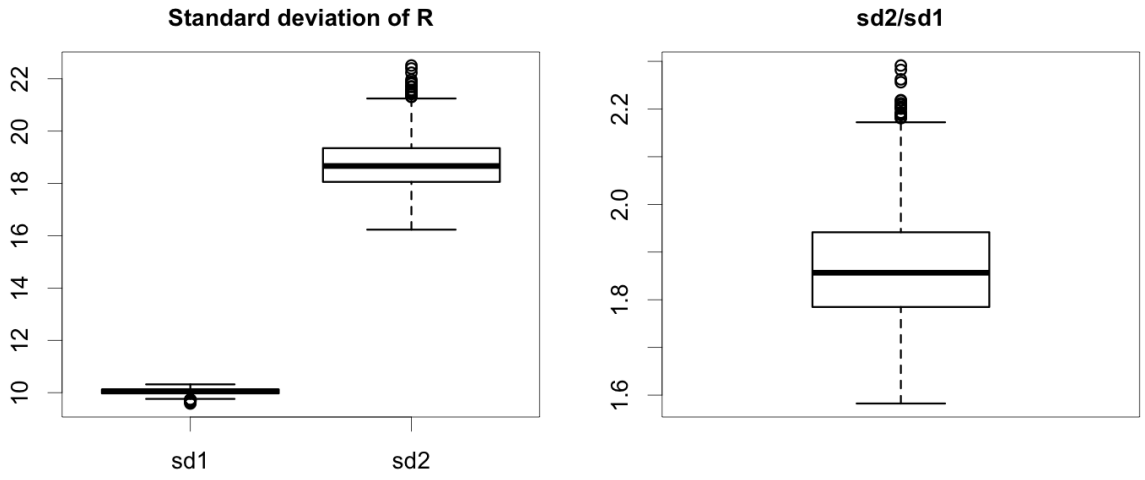

Figure 3: Boxplots of the standard deviation of R before and after adding 50 more observations to one sample, as well as their ratio (after/before), from 1,000 simulation runs.

We see that the standard deviation of R after having the additional 50 observations is on average about 1.85 times as large as that for before, which is much larger than one would expect because the squared root of the ratio of the sample sizes is only $\sqrt{150/100} \approx 1.22$. We see from Figure 2 that the difference between R and $\mathbf{E}(R)$ by having the additional observations is not increased as much as the standard deviation in general, resulting in the decrease in the z-score, $(\mathbf{E}(R) - R)/sd(R)$ (see Figure 4).

Therefore, the decrease of power in scenario 2 is mainly due to the variance boosting problem under unequal sample sizes. To be more exact, from Friedman and Rafsky [1979]

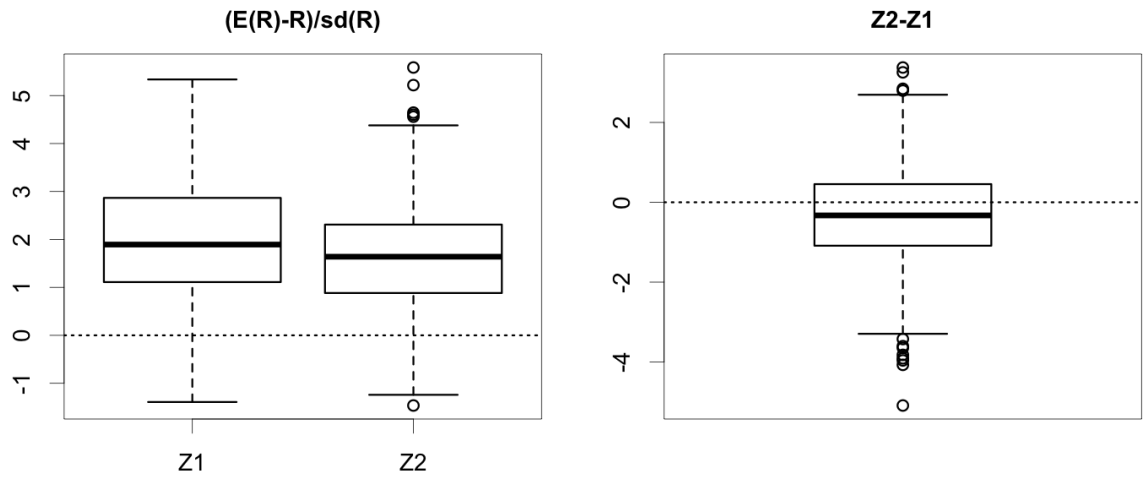

Figure 4: Boxplots of the z-score, $(\mathbf{E}(R) - R)/sd(R)$, before and after adding 50 more observations, as well as their difference (after − before), from 1,000 simulation runs. The horizontal dashed line is at level 0.

and Chen and Friedman [2016], we have the variance of R under the permutation null distribution be

$$\mathbf{Var}(R) = \frac{mn(m-1)(n-1)}{N(N-1)(N-2)(N-3)} \times \left(4|G| + \frac{(n-m)^2 - (N-2)}{(m-1)(n-1)} \left(\sum_{i=1}^{N} |G_i|^2 - \frac{4|G|^2}{N}\right) - \frac{8}{N(N-1)}|G|^2\right), \tag{1}$$

where G_i the subgraph in G that consists of all edge(s) that connect to node i. So $|G_i|$ is the degree of node i in G.

From Cauchy-Schwarz inequality, we know that

$$\sum_{i=1}^{N} |G_i|^2 \geq \frac{4|G|^2}{N},$$

and the equality holds only when $|G_i|$'s are equal for all i's. We call a graph to be *flat* if $\sum_{i=1}^{N} |G_i|^2 - \frac{4|G|^2}{N} = o(N)$, i.e., the degrees of the nodes are similar to each other. When a graph is not flat, such as a k-MST, we see from the expression of $\mathbf{Var}(R)$ that

$$\frac{(n-m)^2}{(m-1)(n-1)} \left(\sum_{i=1}^{N} |G_i|^2 - \frac{4|G|^2}{N}\right) \tag{2}$$

contributes a major portion of the variance when $|m-n| = O(N)$.

Figure 5 plots $\sum_{i=1}^{N} |G_i|^2 - \frac{4|G|^2}{N}$ and $|G|$ for k-MSTs constructed on Euclidean distance in a typical run under scenario 2. We see that $\sum_{i=1}^{N} |G_i|^2 - \frac{4|G|^2}{N}$ is much larger than $|G|$, especially for large k. In scenario 2, $\frac{(n-m)^2}{(m-1)(n-1)} \approx 0.5$. When $k > 3$, the variance of R under scenario 2 is considerably larger than the corresponding case under scenario 1.

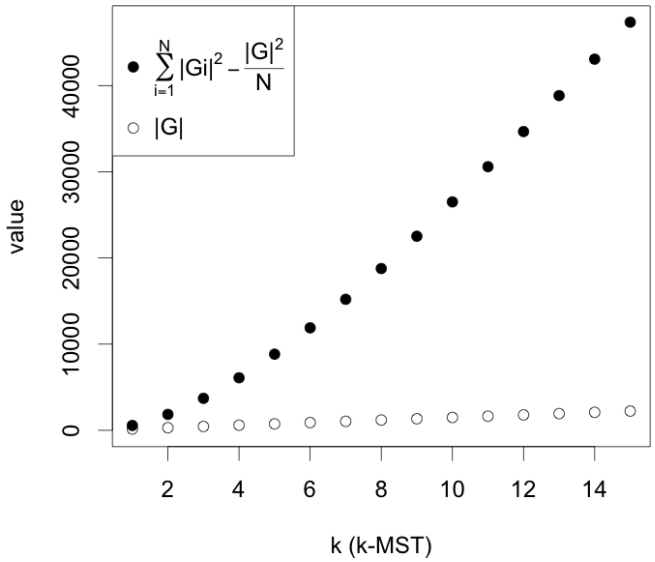

Figure 5: The values of $\sum_{i=1}^{N} |G_i|^2 - \frac{4|G|^2}{N}$ and $|G|$ for a k-MST constructed on Euclidean distance in a typical run under scenario 2.

To avoid the variance boosting problem, we may choose flat graphs, such as k-MDPs, when the sample sizes of the two samples are different. However, this restricts our choices of the similarity graph. Instead, we would rather have a test that does not have the variance boosting problem and works for general graphs.

3 Weighted edge-count test

In this section, we seek to construct a test statistic that captures the signal in a similar way to the edge-count test while not having the variance boosting problem under unequal sample sizes for any graph.

The rationale of the edge-count test is that, under some alternatives, the observations

from the same sample tend to connect within themselves, so the number of *between-sample* edges tends to be *less* than its null expectation, or the number of *within-sample* edges tends to be *more* than its null expectation. Let R_1 be the number of edges that connect within the m observations in sample 1 and R_2 be that for sample 2. Then the test statistic for the edge-count test is equivalent to $R_1 + R_2$. This test statistic is not ideal under unequal sample sizes as it treats the two samples in the same way. However, forming an edge within the sample with a smaller sample size is harder than that for the other sample. We thus propose to weight the within-sample edge counts by the reciprocal of their corresponding sample sizes. In particular, we consider the following two test statistics:

$$R_w = qR_1 + pR_2, \quad p = \frac{m}{N},\ q = 1 - p, \tag{3}$$

$$\tilde{R}_w = \tilde{q}R_1 + \tilde{p}R_2, \quad \tilde{p} = \frac{m-1}{N-2},\ \tilde{q} = 1 - \tilde{p}. \tag{4}$$

Both of their variances are well controlled no matter how different m and n are, and $\tilde{R}_w$ has the smallest variance among all tests of the form $aR_1 + (1-a)R_2$.

Theorem 3.1. *When $m, n \geq 2$, both $\mathbf{Var}(R_w)$ and $\mathbf{Var}(\tilde{R}_w)$ are bounded by*

$$\frac{mn(m-1)(n-1)}{N(N-1)(N-2)(N-3)}|G|. \tag{5}$$

Proof. Following the expressions of $\mathbf{Var}(R_1)$, $\mathbf{Var}(R_2)$, $\mathbf{Cov}(R_1, R_2)$ in the proof of Theorem 3.2 in Chen and Friedman [2016], we have

$$\begin{aligned}\mathbf{Var}(R_w) &= q^2\,\mathbf{Var}(R_1) + p^2\,\mathbf{Var}(R_2) + 2\,p\,q\,\mathbf{Cov}(R_1, R_2) \\ &= \frac{mn(m-1)(n-1)}{N(N-1)(N-2)(N-3)}\times \\ &\quad \left(|G| - \frac{mnN - 2m^2 - 2n^2 + 2mn}{N^2(m-1)(n-1)}\left(\sum_{i=1}^{N}|G_i|^2 - \frac{4|G|^2}{N}\right) - \frac{2}{N(N-1)}|G|^2\right).\end{aligned} \tag{6}$$

Since $\sum_{i=1}^{N}|G_i|^2 \geq \frac{4|G|^2}{N}$ from the Cauchy-Schwarz inequality, and $mnN - 2m^2 - 2n^2 + 2mn = m^2(n-2) + n^2(m-2) + 2mn \geq 0$ when $m, n \geq 2$, we have $\mathbf{Var}(R_w)$ bounded by (5).

For $\mathbf{Var}(\tilde{R}_w)$, we have

$$\begin{aligned}\mathbf{Var}(\tilde{R}_w) &= \tilde{q}^2\,\mathbf{Var}(R_1) + \tilde{p}^2\,\mathbf{Var}(R_2) + 2\,\tilde{p}\,\tilde{q}\,\mathbf{Cov}(R_1, R_2) \\ &= \frac{mn(m-1)(n-1)}{N(N-1)(N-2)(N-3)}\times \\ &\quad \left(|G| - \frac{1}{N-2}\left(\sum_{i=1}^{N}|G_i|^2 - \frac{4|G|^2}{N}\right) - \frac{2}{N(N-1)}|G|^2\right),\end{aligned} \tag{7}$$

and the result follows straightforwardly. □

Theorem 3.2. *For all test statistic of the form $aR_1 + bR_2,\ a + b = 1$, we have*

$$\mathbf{Var}(aR_1 + bR_2) \geq \mathbf{Var}(\tilde{R}_w).$$

Proof. Since

$$\begin{aligned}&\mathbf{Var}(aR_1 + bR_2) = a^2\,\mathbf{Var}(R_1) + b^2\,\mathbf{Var}(R_2) + 2\,a\,b\,\mathbf{Cov}(R_1, R_2) \\ &= \frac{mn(m-1)(n-1)}{N(N-1)(N-2)(N-3)}\times \\ &\quad \left(|G| + \left(a^2\frac{m-2}{n-1} + b^2\frac{n-2}{m-1} - 2ab\right)\left(\sum_{i=1}^{N}|G_i|^2 - \frac{4|G|^2}{N}\right) - \frac{2}{N(N-1)}|G|^2\right),\end{aligned}$$

and

$$\begin{aligned}g(a) &= a^2\frac{m-2}{n-1} + b^2\frac{n-2}{m-1} - 2ab \\ &= a^2\frac{(N-2)(N-3)}{(n-1)(m-1)} - 2a\frac{N-3}{m-1} + \frac{n-2}{m-1}\end{aligned}$$

is minimized when $a = (n-1)/(N-2)$, the result follows. □

Remark 3.3. *When $m = n$, we have $p = q = 0.5$, $\tilde{p} = \tilde{q} = 0.5$, then $R_w = \tilde{R}_w = (R_1 + R_2)/2 = (|G| - R)/2$. So the tests based on R_w, $\tilde{R}_w$ and R are all equivalent under the balanced design.*

Remark 3.4. *For k-MDP, when N is even, every node has degree k, so*

$$2R_1 + R = km,\ \textit{and}\ 2R_2 + R = kn.$$

Then

$$R_w = qR_1 + pR_2 = \frac{n}{N}\left(\frac{km-R}{2}\right) + \frac{m}{N}\left(\frac{kn-R}{2}\right) = \frac{kmn}{N} - \frac{R}{2},$$
$$\tilde{R}_w = \tilde{q}R_1 + \tilde{p}R_2 = \frac{n-1}{N-2}\left(\frac{km-R}{2}\right) + \frac{m-1}{N-2}\left(\frac{kn-R}{2}\right) = \frac{2kmn-kN}{2(N-2)} - \frac{R}{2}.$$

So for k-MDP, when N is even, the tests based on R_w, $\tilde{R}_w$ and R are all equivalent.

When N is odd, the degrees of the nodes are not exactly the same. However, there are at least $N-k$ nodes with degree k and at most k nodes with degree less than k, so the tests based on R_w, $\tilde{R}_w$, and R are all very similar. The same argument holds for any flat graphs. This complies with the earlier observation that the variance boosting problem does not exist for flat graphs.

Remark 3.5. *Asymptotically, when $m, n = O(N)$ and $N \to \infty$, the tests based on R_w and $\tilde{R}_w$ are the same. For finite samples, it turns out that, even though $\mathbf{Var}(\tilde{R}_w)$ is slightly smaller than $\mathbf{Var}(R_w)$, the power of the test based on R_w is slightly higher than that on $\tilde{R}_w$ under locational alternatives. See Section 4.2 for more details of their comparisons.*

4 Power analysis

We first check if the weighted edge-count test solves the variance boosting problem. We compare it to the edge-count test and the generalized edge-count test. For moderate sample sizes, the power of the tests based on R_w and $\tilde{R}_w$ are very similar, so we only include in the comparison the test on R_w (Section 4.1). We then explore the power differences between R_w and $\tilde{R}_w$ for small m and n (Section 4.2).

4.1 A comparison to existing tests

Consider the illustration example in Section 2, we added in the comparison the weighted edge-count test (R_w) and the generalized edge-count test (S) (Figure 6).

We see that the weighted edge-count test (blue triangles) has higher power than the edge-count test when the sample sizes are different for all k-MSTs, $k = 1, \ldots, 15$. It also has higher power than the corresponding smaller sample size scenario ($m = n = 50$). Hence,

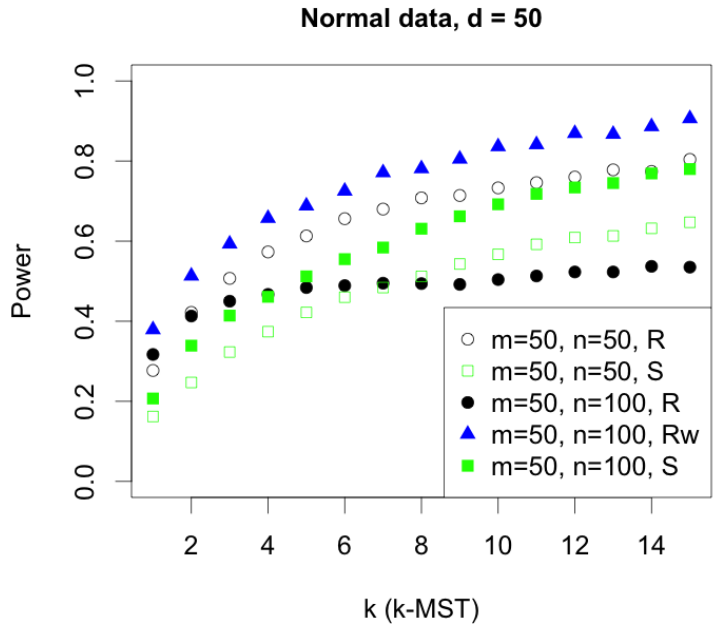

Figure 6: The fraction of trials (out of 1,000) that the test rejected the null hypothesis at 0.05 significance level.

by controlling the variance, the weighted edge-count test does solve the variance boosting problem in the edge-count test.

Comparing the weighted edge-count test to the generalized edge-count test proposed by Chen and Friedman [2016], we see that the weighted edge-count test does have higher power than the generalized edge-count test under locational alternatives, and the weighted edge-count test under scenario 2 ($m = 50, n = 100$) is the only test that has higher power than the edge-count test under scenario 1 ($m = n = 50$) for all k-MSTs.

We checked the performances of the tests under different dimensions (Figure 7). The mean differences are chosen so that the tests have moderate power. We see the same pattern: Under locational alternatives, the weighted edge-count test is more powerful than the edge-count test and the generalized edge-count test when sample sizes are different, and the weighted edge-count test is the only test in scenario 2 that has higher power than the edge-count test in scenario 1 for all k-MSTs.

We also compared all tests for t-distributed data to check how the tests behave when the tail of the distribution is heavier than the normal distribution. The distributions are

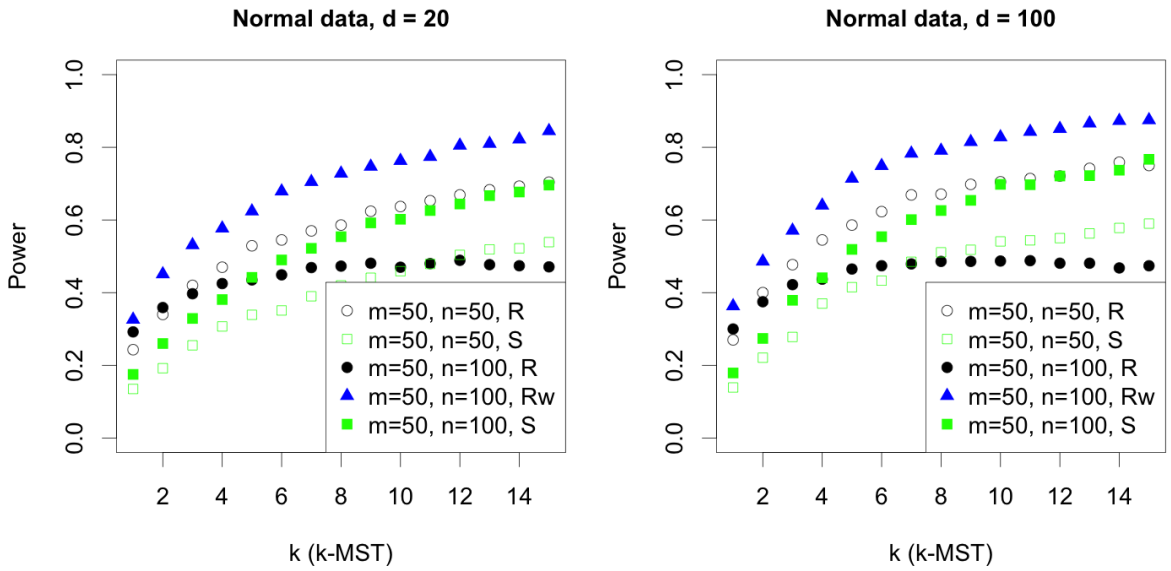

Figure 7: The fraction of trials (out of 1,000) that the test rejected the null hypothesis at 0.05 significance level. The mean differences are of Euclidean distance 1 for $d = 20$ and 1.5 for $d = 100$.

products of independent t distributions. The two distributions differ in the mean, and for each dimension, the mean difference is set to be the same as that under the normal case. The results for t_{10} are shown in Figure 8, and those for t_5 are shown in Figure 9. We see that, overall, all tests have lower power for t-distributed data than for normal data, which indicates that the power of all these tests decreases when the tail of the distribution becomes heavier. Among the three tests, the same pattern retains that the weighted edge-count test outperforms both other tests.

4.2 A comparison between the two weighted edge-count test statistics

According to the definitions of R_w and $\tilde{R}_w$, the tests based on R_w and $\tilde{R}_w$ are very similar for large m and n. We here check the power of them under small m and n. We consider the testing of two samples with one sample $m = 10$ observations from $\mathcal{N}(\mathbf{0}, I_d)$ and the other sample $n = 20$ observations from $\mathcal{N}(\boldsymbol{\mu}, I_d)$, $\|\boldsymbol{\mu}\|_2 = 2,\ d = 50$. The p-values are calculated

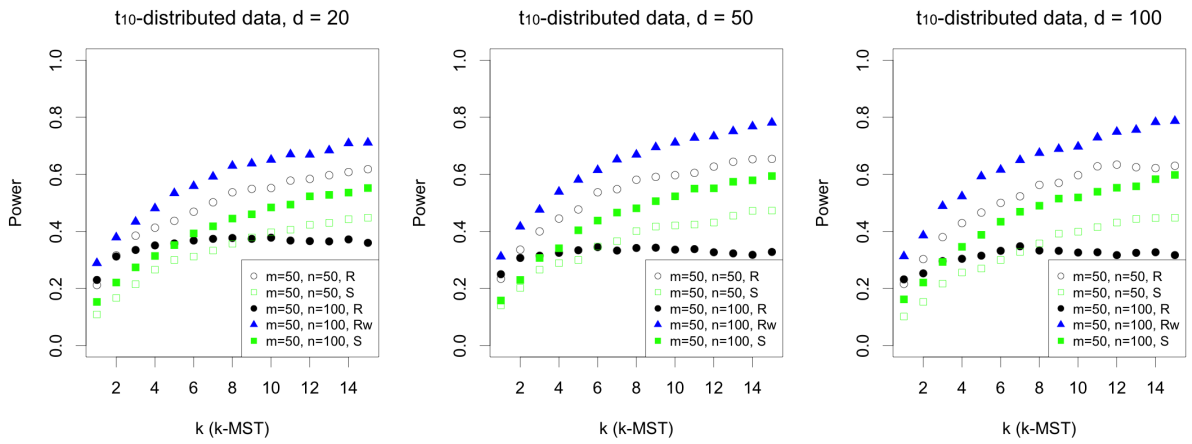

Figure 8: The fraction of trials (out of 1,000) that the test rejected the null hypothesis at 0.05 significance level for t_{10}-distributed data.

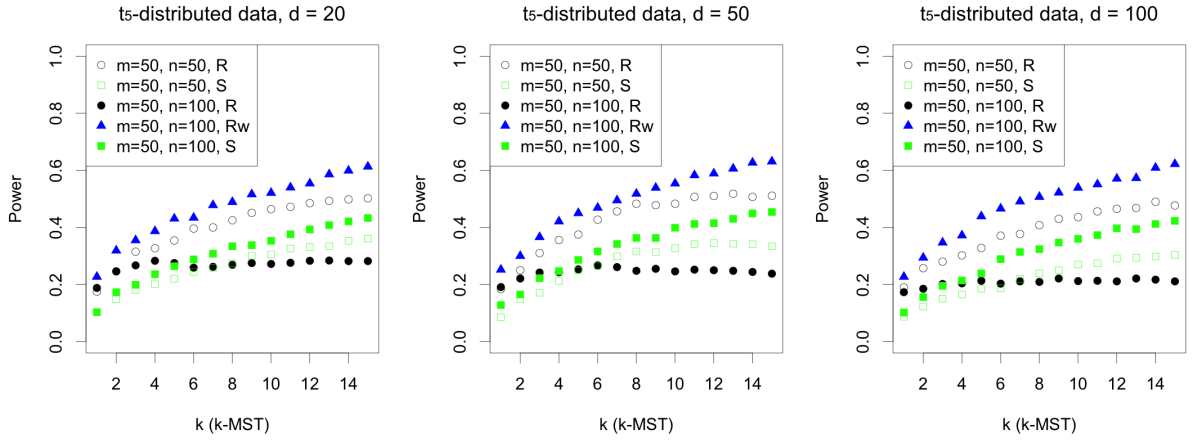

Figure 9: The fraction of trials (out of 1,000) that the test rejected the null hypothesis at 0.05 significance level for t_5-distributed data.

based on 1,000 permutations and the fraction of trials (out of 100) that the test rejected the null hypothesis at 0.05 significance level is listed in Table 1.

Table 1: The fraction of trials (out of 100) that the test rejected the null hypothesis of equal distribution at 0.05 significance level.

	1-MST	2-MST	3-MST	4-MST	5-MST	6-MST	7-MST	8-MST	9-MST
R_w	0.42	0.46	0.46	0.52	0.57	0.57	0.55	0.50	0.52
$\tilde{R}_w$	0.36	0.41	0.41	0.50	0.54	0.54	0.53	0.48	0.48

We see that the test based on R_w has slightly higher power than that based on $\tilde{R}_w$ for all k-MSTs. To check in a more detailed level, we calculate the difference in p-values (p-value of the test based on R_w minus the p-value of the test based on $\tilde{R}_w$) for each trial and the boxplots of the differences for each k-MST, $k = 1, \ldots, 9$, are shown in Figure 10. It is clear that the test based on R_w in general has a smaller p-value than the test based on $\tilde{R}_w$.

According to the above comparison and the simpler form of R_w, we recommend to use R_w in practice.

5 Asymptotics

When the sample size is small, we can obtain the p-value of the test directly through permutations. When the sample size is large, this can be very time consuming. In the following, we study the asymptotic distribution of R_w under the *usual limiting regime*, which is defined as $|G|, m, n \to \infty, m/N \to \lambda \in (0, 1)$. We show that, in the usual limiting regime, R_w, normalized by its mean and standard deviation, approaches the standard normal distribution as $N \to \infty$ under some mild conditions on the similarity graph G. We then check how well the asymptotic null distribution works in approximating p-values for finite samples.

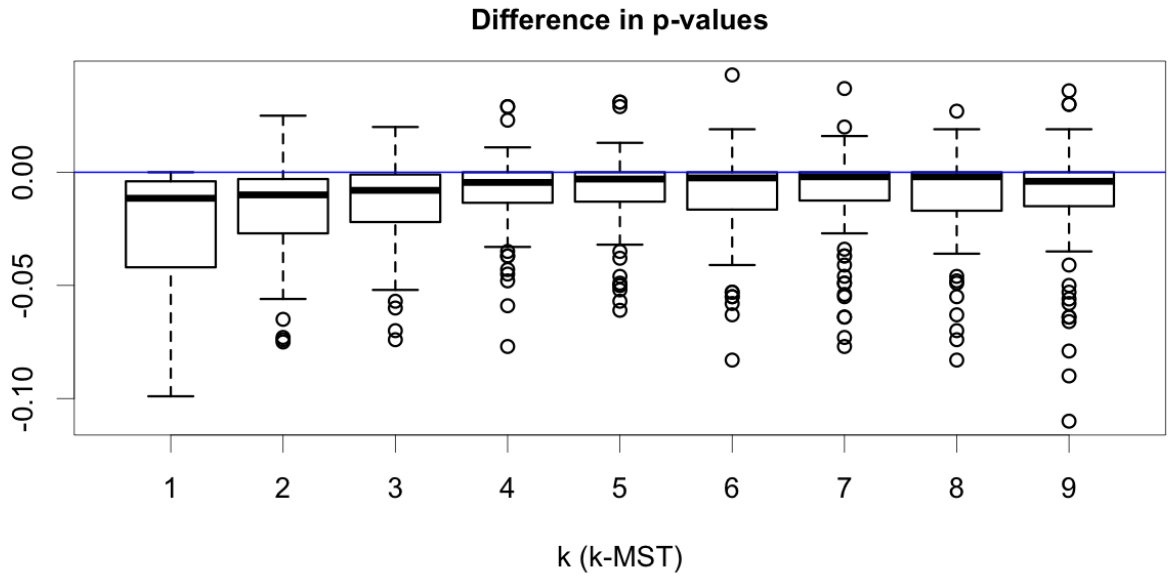

Figure 10: Boxplots of the difference between the p-values of the tests based on R_w and $\tilde{R}_w$ (the p-value of the test based on R_w minus the p-value of the test based on $\tilde{R}_w$) for k-MST, $k = 1, \ldots, 9$. The horizontal line is at level 0.

5.1 Asymptotic null distribution

Before stating the theorem, we define two additional terms on the similarity graph G.

$$A_e = \{e\} \cup \{e' \in G : e' \text{ and } e \text{ share a node}\},$$
$$B_e = A_e \cup \{e'' \in G : \exists\ e' \in A_e, \text{ such that } e'' \text{ and } e' \text{ share a node}\}.$$

So A_e is the subgraph in G that consists of all edge(s) that connect to edge e, and B_e is the subgraph in G that consists of all edge(s) that connect to any edge in A_e.

Theorem 5.1. *If* $|G| = O(N^\alpha),\ 1 \le \alpha < 1.5$, $\sum_{e\in G} |A_e||B_e| = o(N^{1.5\alpha})$, *and* $\sum_{e\in G} |A_e|^2 = o(N^{\alpha+0.5})$, *in the usual limiting regime, under the permutation null,*

$$Z_w := \frac{R_w - \mathbf{E}(R_w)}{\sqrt{\mathbf{Var}(R_w)}} \xrightarrow{\mathcal{D}} \mathcal{N}(0,1). \tag{8}$$

The proof for this theorem utilizes Stein's method [Chen and Shao, 2005]. The complete proof is in Appendix A.

Remark 5.2. *This theorem also holds for* $|G| = O(N^\alpha),\ 0.5 < \alpha < 1$, *along with some other conditions on the graph. However, such similarity graphs do not connect most of*

the nodes and are thus not of interest in practice as they missed most of the similarity information among the observations.

Remark 5.3. *The conditions* $\sum_{e\in G}|A_e||B_e| = o(N^{1.5\alpha})$ *and* $\sum_{e\in G}|A_e|^2 = o(N^{\alpha+0.5})$ *ensure that the graph does not have a huge hub or a cluster of small hubs, where a hub is a node with a large degree. If we only concern graphs with* $|G| = O(N)$*, i.e.,* $\alpha = 1$*, then these two conditions degenerate into one condition:* $\sum_{e\in G}|A_e||B_e| = o(N^{1.5})$*. Hence, the conditions in Theorem 5.1 are much more relaxed than the conditions in obtaining the limiting distribution for the generalized edge-count test statistic in Chen and Friedman [2016], which include not only* $|G| = O(N)$ *and* $\sum_{e\in G}|A_e||B_e| = o(N^{1.5})$*, but also* $\sum_{i=1}^{N}|G_i|^2 = O(N)$ *and* $\sum_{i=1}^{N}|G_i|^2 - 4|G|^2/N = O(N)$*.*

Corollary 5.4. *When the graph is a k-MST, where* $k = O(1)$*, based on the Euclidean distance,* $Z_w \xrightarrow{\mathcal{D}} \mathcal{N}(0,1)$ *in the usual limiting regime under the null hypothesis.*

A k-MST, where $k = O(1)$, constructed on the Euclidean distance satisfies all conditions in obtaining the limiting distribution for the generalized edge-count test statistic [Chen and Friedman, 2016]. Hence, it satisfies all conditions for Theorem 5.1, and the result follows.

5.2 Consistency

Theorem 5.5. *For two continuous multivariate distributions, if the graph is a k-MST,* $k = O(1)$*, based on the Euclidean distance, the weighted edge-count test is consistent against all alternatives in the usual limiting regime.*

This theorem can be proved through arguments extended from Henze and Penrose [1999]. The details are in Appendix B.

5.3 Accuracy of the p-value approximation from the asymptotics for finite sample sizes

We here check how large the sample sizes need to be so that the asymptotic p-value based on Theorem 5.1 is a good approximation to the permutation p-value. Figure 11 shows boxplots of the differences of the two p-values (approximated p-value from asymptotic results minus

p-value calculated through 1,000 permutations) from 100 simulation runs under different choices of m, n, d, and the graph. We see from the boxplots that the approximate p-value is very accurate for sample sizes in hundreds. Making the graph slightly denser, or making the ratio of the two sample sizes higher does not affect the accuracy much. Increasing the dimension does not affect the accuracy much either.

$d = 20$:

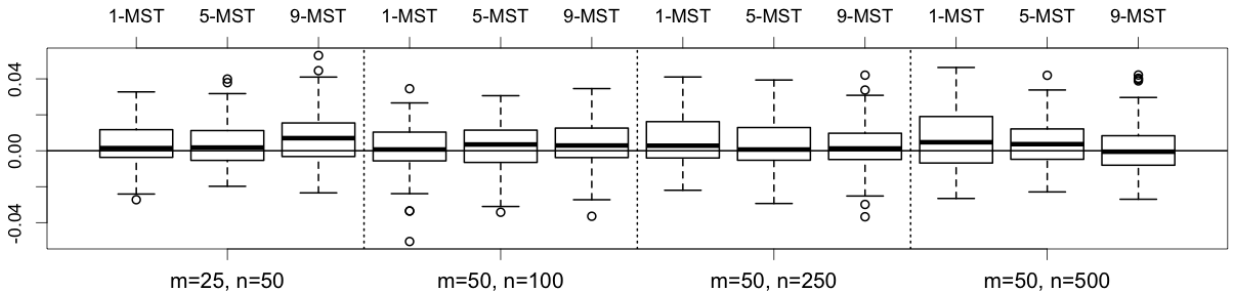

$d = 100$:

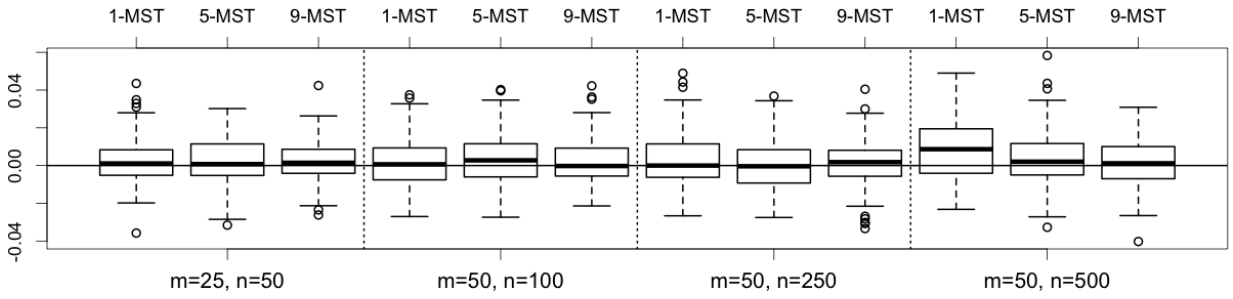

Figure 11: Boxplots of the differences between the p-value based on asymptotic distribution and the p-value calculated directly from 1,000 permutations (100 simulation runs for each setting. $F_1 = F_2 = \mathcal{N}(\mathbf{0}, I_d)$)).

6 A real data example

The MIT Media Laboratory conducted a study following 106 subjects, students and staffs in an institute, who used mobile phones with pre-installed software that can record call logs. The study lasted from July 2004 to June 2005 [Eagle et al., 2009]. Given the richness of this dataset, many problems can be studied. One question of interest is whether phone call patterns on weekdays are different from those on weekends. The phone calls on weekdays

and weekends can be viewed as representations of professional relationship and personal relationship, respectively.

We bin the phone calls by day and, for each day, construct a directed phone-call network with the 106 subjects as nodes and a directed edge pointing from person i to person j if person i made at least one call to person j on that day. We encode the directed network of each day by an adjacency matrix, with 1 for element $[i, j]$ if there is a directed edge pointing from subject i to subject j, and 0 if otherwise.

In this period, there was no call among the subjects on 9.6% of the days in weekdays and 9.3% of the days in weekends. We remove these days and end up with 214 days in weekdays and 85 days in weekends. Let $A_1, \ldots, A_{214}$ be the adjacency matrices on the 214 weekdays and $A_{215}, \ldots, A_{299}$ be that on the 85 days in weekends. We consider two distance measures defined as:

(1) the number of different entries: $d(A_i, A_j) = \|A_i - A_j\|_F^2$, where $\|\cdot\|_F$ is the Frobenius norm of a matrix,

(2) the number of different entries, normalized by the geometric mean of the total edges in each of the two days: $d(A_i, A_j) = \frac{\|A_i - A_j\|_F^2}{\|A_i\|_F \|A_j\|_F}$.

Figure 12 shows the p-values of the weighted edge-count test, the edge-count test, and the generalized edge-count test, on k-MSTs constructed on each of the two distances. We see that, except for the small k's, the weighted edge-count test rejects the null hypothesis of equal distribution at 0.05 significance level under both distances, while the edge-count test rejects none of them. The generalized edge-count test rejects most of the scenarios. Under distance 1, the p-values of the generalized edge-count test scatter around 0.05 for $k = 8, \ldots, 20$.

Since all tests have higher power when the graph is slightly denser, the weighted edge-count test rejects the null hypothesis of equal distribution under both distances, the edge-count test does not reject the null under either distance, and the generalized edge-count test more or less rejects the null under both distances.

In the following, we study several representative cases in more details: 9-MST and 15-MST under distance 1 and 9-MST under distance 2. Table 2 lists the summary statistics

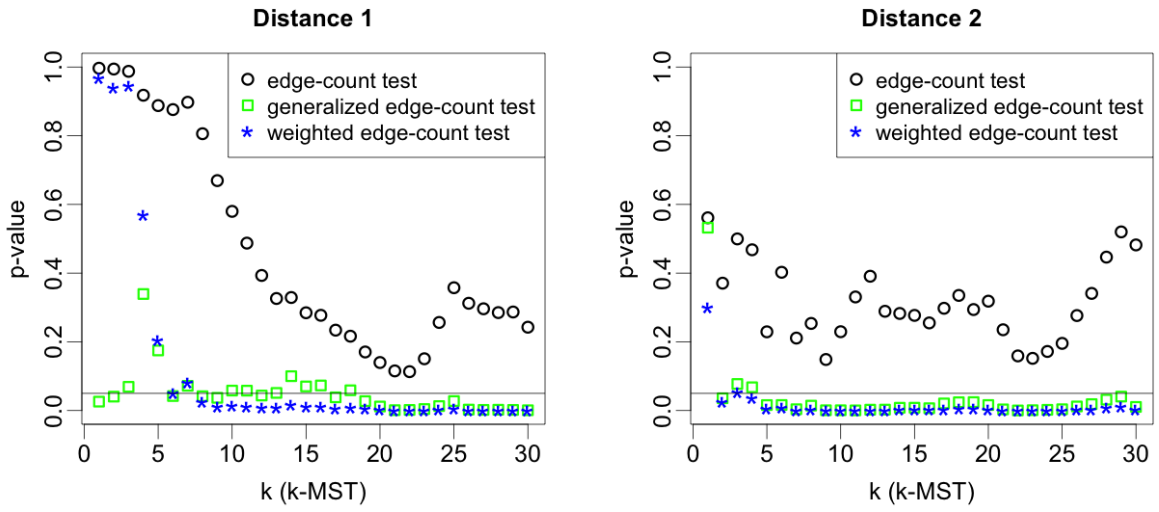

Figure 12: The p-values of the edge-count tests (circles), the generalized edge-count test (squares), and the weighted edge-count test (stars) on k-MSTs constructed on distance 1 and distance 2, respectively. The horizontal line is of level 0.05.

of the tests, in particular, the values of R, R_1, R_2, $(R_1 + R_2)/2$ and R_w, as well as their expectations (mean), standard deviations (SD), and z-scores ((value-mean)/SD). The test based on $(R_1+R_2)/2 = (|G|-R)/2$ is equivalent to that based on R. We include $(R_1+R_2)/2$ in the table to make the comparison between the (unweighted) edge-count test and the weighted edge-count test easier.

First, we take a close look at results on 9-MST based on distance 1. There are less-than-expectation within-sample edges for the weekday sample and more-than-expectation within-sample edges for the weekend sample. The sample size for the weekday sample is about 2.5 times as large as the the sample size for the weekend sample. If we simply add the within-sample edges from the two samples, the number of total within-sample edges is less than its null expectation, falsely indicating that the observations are less likely to form edges within the samples, which leads to the conclusion that the two samples are well mixed and the null hypotheses is not rejected. This is what the (unweighted) edge-count test does.

On the other hand, even though $\mathbf{E}(R_1) - R_1 = 98.03$ is larger than $R_2 - \mathbf{E}(R_2) = 69.08$, it is still more likely to observe values that are more extreme than R_1 than that for R_2 if

Table 2: Summary statistics for 9-MST and 15-MST under distance 1 and 9-MST under distance 2. Sample 1: phone-call networks on weekdays; sample 2: phone-call networks on weekends.

9-MST, distance 1

	Value	Mean	Value − Mean	SD	z-score
R	1124	1095.05	28.95	66.09	0.438
R_1	1274	1372.03	-98.03	104.95	-0.934
R_2	284	214.92	69.08	42.30	1.633
$\frac{R_1+R_2}{2}$	779	793.47	-14.47	33.04	-0.438
R_w	565.44	543.86	21.58	9.50	2.272

15-MST, distance 1

	Value	Mean	Value − Mean	SD	z-score
R	1770	1825.08	-55.08	96.77	-0.569
R_1	2316	2286.72	29.28	155.54	0.188
R_2	384	358.19	25.81	62.26	0.414
$\frac{R_1+R_2}{2}$	1350	1322.46	27.54	48.38	0.569
R_w	933.23	906.44	26.79	11.62	2.305

9-MST, distance 2

	Value	Mean	Value − Mean	SD	z-score
R	1055	1095.05	-40.05	38.41	-1.043
R_1	1354	1372.03	-18.03	54.99	-0.327
R_2	273	214.92	58.08	23.57	2.465
$\frac{R_1+R_2}{2}$	813.5	793.47	20.03	19.21	1.043
R_w	580.31	543.86	36.44	10.04	3.629

we take into account the sample sizes. The weighted test statistic correctly summarizes the signals provided by both samples and results in a larger than expectation weighted within-sample edges. Together with the variance minimizing effect in the weighted edge-count test, the test statistic is significantly enough to reject the null hypothesis.

A similar argument holds for the results on 9-MST based on distance 2. However, in this case, $\mathbf{E}(R_1) - R_1$ is smaller than $R_2 - \mathbf{E}(R_2)$, so the edge-count test also concludes that there are more-than-expectation within-sample edges, while the test is not done effectively that the null is not rejected.

The case of 15-MST based on distance 1 is slightly different from the other two cases that the within-sample edges for both samples are larger than their null expectations. However, the total difference by plain addition is not significant as it is only about half of its corresponding standard deviation. The variance minimizing effect of the weighted edge-count test is well reflected here that R_w is of a similar amount away from its null expectation compared to the unweighted version ($(R_1 + R_2)/2$)), but the standard deviation of R_w is much smaller than that for $(R_1 + R_2)/2$, vastly improving the power to detect the signal. In this case, since both samples are more likely to connect within themselves, the alternative falls in the area that the weighted edge-count test is the most effective. We see that the generalized edge-count test is not powerful enough for this case, and only the weighted edge-count test rejects the null hypothesis at 0.05 significance level among the three tests.

7 Relation between the weighted edge-count test and the generalized edge-count test

The test statistic of the generalized edge-count test is

$$S = (R_1 - \mathbf{E}(R_1), R_2 - \mathbf{E}(R_2))\Sigma_R^{-1} \begin{pmatrix} R_1 - \mathbf{E}(R_1) \\ R_2 - \mathbf{E}(R_2) \end{pmatrix}, \tag{9}$$

where Σ_R is the covariance matrix of the vector $(R_1, R_2)'$.

According to Remark 3.4 in Chen and Friedman [2016], when $\sum_{i=1}^{N} |G_i|^2 - \frac{4|G|^2}{N} =$

$O(|G|)$, which is commonly achieved for k-MST, $k = O(1)$,

$$\lim_{N\to\infty} \frac{\Sigma_R}{|G|} = p^2q^2 \begin{pmatrix} 1+rp/q & 1-r \\ 1-r & 1+rq/p \end{pmatrix}, \tag{10}$$

where $r = \lim_{N\to\infty} \sum_{i=1}^{N} (|G_i|^2 - 4|G|^2/N)/|G|$.

Since

$$\begin{pmatrix} 1+rp/q & 1-r \\ 1-r & 1+rq/p \end{pmatrix} = \begin{pmatrix} p & 1 \\ -q & 1 \end{pmatrix} \begin{pmatrix} \frac{r}{pq} & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} p & -q \\ 1 & 1 \end{pmatrix},$$

we have

$$\begin{pmatrix} 1+rp/q & 1-r \\ 1-r & 1+rq/p \end{pmatrix}^{-1} = \begin{pmatrix} 1 & q \\ -1 & p \end{pmatrix} \begin{pmatrix} \frac{pq}{r} & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & -1 \\ q & p \end{pmatrix},$$

and

$$\lim_{N\to\infty} S = \frac{(qR_1 + pR_2 - (q\mathbf{E}(R_1) + p\mathbf{E}(R_2)))^2}{p^2q^2|G|} + \frac{(R_1 - R_2 - (\mathbf{E}(R_1) - \mathbf{E}(R_2)))^2}{rpq|G|}.$$

We see that the first part of the summation corresponds to the weighted edge-count test. As discussed in Chen and Friedman [2016], under alternative hypothesis, it can be that (i) both samples are more likely to connect within themselves; or (ii) one sample is more likely to connect within themselves while the other sample is less likely to connect within themselves. The two parts in the summation correspond to the two scenarios. So the generalized edge-count test works under a wider range of alternatives than the weighted edge-count test does, while if the alternative is of scenario (i), the weighted edge-count test is slightly more powerful.

Since no alternative would lead to the scenario that both samples are less likely to connect within themselves when testing two randomly drawn samples, to make the generalized edge-count test slightly more powerful for both scenarios, we can use the test statistic $\max(Z_w, |Z_{\text{diff}}|)$ with

$$Z_w = \frac{qR_1 + pR_2 - (q\mathbf{E}(R_1) + p\mathbf{E}(R_2))}{pq\sqrt{|G|}},$$
$$Z_{\text{diff}} = \frac{R_1 - R_2 - (\mathbf{E}(R_1) - \mathbf{E}(R_2))}{\sqrt{rpq|G|}}.$$

It can be shown that $(Z_w, Z_{\text{diff}})'$ is asymptotically bivariate Gaussian distributed and Z_w and Z_{diff} are asymptotically independent, so the asymptotic critical value for this variant generalized edge-count test can be easily determined.

8 Conclusion

We propose a new two-sample test that utilizes the similarity information among the observations. In particular, the test is based on a similarity graph constructed on the pooled observations. Thus, the test can be applied to multivariate data and non-Euclidean data as long as an informative similarity measure on the sample space can be defined. The classic test of this type, the edge-count test, has an issue when the sample sizes are different. The new test solves this problem by giving weights based on sample sizes to the different components of the edge-count test statistic. This weighted edge-count test exhibits substantial power gains in simulation studies.

The weighted edge-count test statistic, standardized by its mean and standard deviation, is shown to converge to the standard normal distribution under some mild conditions on the similarity graph. The approximated p-value based on the asymptotic results is reasonably accurate to the permutation p-value for sample sizes in hundreds, facilitating the application of the test to large datasets.

The weighted edge-count test and the generalized edge-count test in Chen and Friedman [2016] can be used in a complementary way. The generalized edge-count test works for a wider range of alternatives under practical sample sizes compared to the weighted edge-count test. However, under locational alternatives, the power of the weighted edge-count test is higher than that for the generalized edge-count test. The choice of whether to use the weighted edge-count test or the generalized edge-count test can be done in a similar manner to the choice between the Hotelling T^2 test and the generalized likelihood ratio test not assuming equal covariance matrix.

Acknowledgements

Hao Chen is supported in part by NSF award DMS-1513653.

References

Hao Chen and Jerome H Friedman. A new graph-based two-sample test for multivariate and object data. *Journal of the American statistical association*, 2016.

Hao Chen and Nancy R. Zhang. Graph-based tests for two-sample comparisons of categorical data. *Statistica Sinica*, 23:1479–1503, 2013.

Louis H.Y. Chen and Qi-Man Shao. Stein's method for normal approximation. *An introduction to Stein's method*, Lecture Notes Series No. 4, Institute for Mathematical Sciences, National University of Singapore, Singapore University Press and World Scientific:1–59, 2005.

Alejandro de la Sierra, Julián Segura, José R Banegas, Manuel Gorostidi, J Juan, Pedro Armario, Anna Oliveras, and Luis M Ruilope. Clinical features of 8295 patients with resistant hypertension classified on the basis of ambulatory blood pressure monitoring. *Hypertension*, 57(5):898–902, 2011.

Nathan Eagle, Alex Sandy Pentland, and David Lazer. Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences*, 106 (36):15274–15278, 2009.

Keith A Feigenson, Michael A Gara, Matthew W Roché, and Steven M Silverstein. Is disorganization a feature of schizophrenia or a modifying influence: Evidence of covariation of perceptual and cognitive organization in a non-patient sample. *Psychiatry research*, 217(1):1–8, 2014.

Jerome H. Friedman and Lawrence C. Rafsky. Multivariate generalizations of the Wald-Wolfowitz and Smirnov two-sample tests. *The Annals of Statistics*, 7(4):697–717, 1979.

Jean Dickinson Gibbons and Subhabrata Chakraborti. *Nonparametric statistical inference*. Springer, 2011.

Norbert Henze. A multivariate two-sample test based on the number of nearest neighbor type coincidences. *The Annals of Statistics*, 16(2):772–783, 1988.

Norbert Henze and Mathew D. Penrose. On the multivariate runs test. *Annals of statistics*, 27(1):290–298, 1999.

Gueorgi Kossinets and Duncan J Watts. Empirical analysis of an evolving social network. *Science*, 311(5757):88–90, 2006.

Jen-Fue Maa, Dennis K. Pearl, and Robert Bartoszyński. Reducing multidimensional two-sample data to one-dimensional interpoint comparisons. *The Annals of Statistics*, 24(3): 1069–1074, 1996.

Paul R. Rosenbaum. An exact distribution-free test comparing two multivariate distributions based on adjacency. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(4):515–530, 2005.

Mark F. Schilling. Multivariate two-sample tests based on nearest neighbors. *Journal of the American Statistical Association*, 81(395):799–806, 1986.

A Proof to Theorem 5.1

The proof of Theorem 5.1 relies on Stein's method. Consider sums of the form $W = \sum_{i\in\mathcal{J}} \xi_i$, where $\mathcal{J}$ is an index set and ξ are random variables with $\mathbf{E}\xi_i = 0$, and $\mathbf{E}(W^2) = 1$. The following assumption restricts the dependence between $\{\xi_i : \ i \in \mathcal{J}\}$.

Assumption A.1. *[Chen and Shao, 2005, p. 17] For each $i \in \mathcal{J}$ there exists $S_i \subset T_i \subset \mathcal{J}$ such that ξ_i is independent of $\xi_{S_i^c}$ and ξ_{S_i} is independent of $\xi_{T_i^c}$.*

We will use the following theorem in proving Theorem 5.1.

Theorem A.1. *[Chen and Shao, 2005, Theorem 3.4] Under Assumption A.1, we have*

$$\sup_{h\in Lip(1)} |\mathbf{E}h(W) - \mathbf{E}h(Z)| \leq \delta,$$

where $Lip(1) = \{h : \mathbb{R} \to \mathbb{R}\}$, Z has $\mathcal{N}(0,1)$ distribution and

$$\delta = 2\sum_{i\in\mathcal{J}} (\mathbf{E}|\xi_i\eta_i\theta_i| + |\mathbf{E}(\xi_i\eta_i)|\mathbf{E}|\theta_i|) + \sum_{i\in\mathcal{J}} \mathbf{E}|\xi_i\eta_i^2|$$

with $\eta_i = \sum_{j\in S_i} \xi_j$ and $\theta_i = \sum_{j\in T_i} \xi_j$, where S_i and T_i are defined in Assumption A.1.

To prove Theorem 5.1, we take one step back to study the statistic under the bootstrap null distribution, which is defined as follows: For each observation, we assign it to be from sample 1 with probability m/N, and from sample 2 with probability n/N, independently of other observations. We use g_i to denote the sample assignment for observation i, with $g_i = 1$ if observation i is assigned to be from sample 1 and $g_i = 2$ if otherwise. Let $X = \sum_{i=1}^{N} I(g_i = 1)$ be the number of observations assigned to be from sample 1, where $I(\cdot)$ is the indicator function. Then, the bootstrap null distribution conditioning on $X = m$ is equivalent to the permutation null distribution. We use $\mathbf{P_B}$, $\mathbf{E_B}$, $\mathbf{Var_B}$, $\mathbf{Cov_B}$ to denote the probability, expectation, variance, and covariance under the bootstrap null distribution, respectively. (We here add the subscript $_{\mathbf{P}}$ to denote the corresponding quantities under the permutation null distribution.)

Let $p_N = m/N, q_N = 1 - p_N$, then $\lim_{N\to\infty} p_N = \lambda, \lim_{N\to\infty} q_N = 1 - \lambda$. For any $e = (e_-, e_+) \in G$, let

$$\beta_e = \begin{cases} 1 & \text{if } g_{e_-} = g_{e_+} = 1, \\ 2 & \text{if } g_{e_-} = g_{e_+} = 2, \\ 0 & \text{if } g_{e_-} \neq g_{e_+}. \end{cases}$$

Given that the g_i's are independent under the bootstrap null distribution, we have

$$\begin{aligned} \mathbf{E_B}(R_1) &= \frac{m^2}{N^2}|G|, \\ \mathbf{E_B}(R_2) &= \frac{n^2}{N^2}|G|, \\ \mathbf{Var_B}(R_1) &= \frac{m^2n^2}{N^4}|G| + \frac{m^3n}{N^4}\sum_{i=1}^{N}|G_i|^2, \\ \mathbf{Var_B}(R_2) &= \frac{m^2n^2}{N^4}|G| + \frac{mn^3}{N^4}\sum_{i=1}^{N}|G_i|^2, \\ \mathbf{Cov_B}(R_1, R_2) &= \frac{m^2n^2}{N^4}|G| - \frac{m^2n^2}{N^4}\sum_{i=1}^{N}|G_i|^2. \end{aligned}$$

For $R_w = q_N R_1 + p_N R_2$, we have

$$\begin{aligned} \mathbf{E_B}(R_w) &= \frac{mn}{N^2}|G| := \mu_B, \\ \mathbf{Var_B}(R_w) &= \frac{m^2n^2}{N^4}|G| := \sigma_B^2. \end{aligned}$$

In contrast, we have

$$\begin{aligned}
\mathbf{E_P}(R_w) &= \frac{mn(N-2)}{N^2(N-1)}|G| := \mu_P, \\
\mathbf{Var_P}(R_w) &\frac{mn(m-1)(n-1)}{N(N-1)(N-2)(N-3)} \times \\
&\left(|G| - \frac{mnN - 2m^2 - 2n^2 + 2mn}{N^2(m-1)(n-1)} \left(\sum_{i=1}^{N} |G_i|^2 - \frac{4|G|^2}{N} \right) - \frac{2}{N(N-1)}|G|^2 \right) \\
&:= \sigma_P^2,
\end{aligned}$$

Let

$$Z_w^B = \frac{R_w - \mu_B}{\sigma_B}, \quad Z_X^B = \frac{X - m}{\sqrt{mn/N}}.$$

Under the conditions of Theorem 5.1, as $N \to \infty$, we can prove the following results:

(i) (Z_w^B, Z_X^B) becomes bivariate Gaussian distributed under the bootstrap null.

(ii)

$$\frac{\sigma_P}{\sigma_B} \to 1, \quad \frac{\mu_B - \mu_P}{\sigma_P} \to 0.$$

From (i) and together with $\mathbf{Var_B}(Z_X^B) = 1$, as $N \to \infty$, the conditional distribution of Z_w^B given Z_X^B converges to a Gaussian distribution under the bootstrap null distribution. Since the bootstrap null distribution conditioning on $Z_X^B = 0$ is equivalent to the permutation null distribution, Z_w^B follows a Gaussian distribution under the permutation null distribution as $N \to \infty$. Notice that

$$Z_w = \frac{\sigma_B}{\sigma_P} \left(Z_w^B + \frac{\mu_B - \mu_P}{\sigma_B} \right),$$

together with (ii), we have Z_w converges to a Gaussian distribution under the permutation null distribution as $N \to \infty$.

In the following, we prove results (i) and (ii).

To prove (i), by Cramér-Wold device, we only need to show that $W = a_1 Z_w^B + a_2 Z_X^B$ is asymptotically Gaussian distributed for any combination of a_1 and a_2 such that $\mathbf{Var_B}(W) > 0$.

Let

$$\xi_e = \frac{a_1}{\sigma_B}\left(\frac{n}{N}I_{\beta_e=1} + \frac{m}{N}I_{\beta_e=2} - \frac{mn}{N^2}\right),$$
$$\xi_i = a_2\frac{I_{g_i=1} - m/N}{\sigma_0},\ \sigma_0 = \sqrt{mn/N}.$$

Then $W = \sum_{j\in\mathcal{J}}\xi_j$, where $\mathcal{J} = \{e \in G\} \cup \{1, \ldots, N\}$. Let $a = \max(|a_1|, |a_2|)$, then $|\xi_e| \le a/\sigma_B$ and $|\xi_i| \le a/\sigma_0$.

For $e = (e_-, e_+) \in \mathcal{J}$, let

$$S_e = A_e \cup \{e_-, e_+\},$$
$$T_e = B_e \cup \{\text{nodes in } A_e\}.$$

Then S_e and T_e satisfy Assumption A.1.

For $i \in \{1, \ldots, N\}$, let

$$S_i = \{e \in G_i\} \cup \{i\},$$
$$T_i = \{e \in G_{i,2}\} \cup \{\text{nodes in } G_i\},$$

where $G_{i,2} = \{(j,l) \in G : j \in G_i\}$ is the subgraph is G that consists of all edges that connect to any node in G_i. Then S_i and T_i satisfy Assumption A.1.

For $j \in \mathcal{J}$, let $\eta_j = \sum_{k\in K_j}\xi_k$, $\theta_j = \sum_{k\in L_j}\xi_k$. By Theorem A.1, we have $\sup_{h\in Lip(1)}|\mathbf{E_B}h(W) - \mathbf{E}h(Z)| \le \delta$ for $Z \sim \mathcal{N}(0,1)$, where

$$\begin{aligned}\delta &= \frac{1}{\sqrt{\mathbf{Var_B}(W)}}\left(2\sum_{j\in\mathcal{J}}(\mathbf{E_B}|\xi_j\eta_j\theta_j| + |\mathbf{E_B}(\xi_j\eta_j)|\mathbf{E_B}|\theta_j|) + \sum_{j\in\mathcal{J}}\mathbf{E_B}|\xi_j\eta_j^2|\right)\\ &\le \frac{a^3}{\sqrt{\mathbf{Var_B}(W)}}\left(5\sum_{e\in G}\frac{1}{\sigma_B}\left(\frac{|A_e|}{\sigma_B} + \frac{2}{\sigma_0}\right)\left(\frac{|B_e|}{\sigma_B} + \frac{2|A_e|}{\sigma_0}\right)\right.\\ &\qquad\left. + 5\sum_{i=1}^{N}\frac{1}{\sigma_0}\left(\frac{|G_i|}{\sigma_B} + \frac{1}{\sigma_0}\right)\left(\frac{|G_{i,2}|}{\sigma_B} + \frac{2|G_i|}{\sigma_0}\right)\right).\end{aligned}$$

Since $\sigma_B = O(|G|^{0.5})$ and $\sigma_0 = O(N^{0.5})$, then as long as

$$\sum_{e\in G} |A_e||B_e| = o(|G|^{1.5}), \tag{11}$$

$$\sum_{e\in G} |A_e|^2 = o(|G| \cdot N^{0.5}), \tag{12}$$

$$\sum_{e\in G} |B_e| = o(|G| \cdot N^{0.5}), \tag{13}$$

$$\sum_{e\in G} |A_e| = o(|G|^{0.5} \cdot N), \tag{14}$$

$$\sum_{i=1}^{N} |G_i||G_{i,2}| = o(|G| \cdot N^{0.5}), \tag{15}$$

$$\sum_{i=1}^{N} |G_i|^2 = o(|G|^{0.5} \cdot N), \tag{16}$$

$$\sum_{i=1}^{N} |G_{i,2}| = o(|G|^{0.5} \cdot N), \tag{17}$$

$$\sum_{i=1}^{N} |G_i| = o(N^{1.5}), \tag{18}$$

result (i) follows. We next show that the conditions in Theorem 5.1 ($|G| = O(N^\alpha),\ 1 \le \alpha < 1.5$, $\sum_{e\in G} |A_e||B_e| = o(N^{1.5\alpha})$, $\sum_{e\in G} |A_e|^2 = o(N^{\alpha+0.5})$) are enough to show (11)-(18).

First of all, if we substitute $|G|$ by $O(N^\alpha)$, (11), (12), and (18) follow immediately.

Since $|B_e| \le \sum_{e'\in A_e} |A_{e'}|$, we have $\sum_{e\in G} |B_e| \le \sum_{e\in G} \sum_{e'\in A_e} |A_{e'}|$. For the latter quantity, for any $e^* \in G$, $|A_{e^*}|$ appears $|A_{e^*}|$ times because e^* is the first neighbor of $(|A_{e^*}| - 1)$ edges. Hence, the latter quantity equals $\sum_{e\in G} |A_e|^2$. So (12) ensures (13).

By Cauchy-Schwarz, we have $\sum_{e\in G} |A_e| \le \sqrt{\sum_{e\in G} |A_e|^2 \sum_{e\in G} 1^2} = o(|G| \cdot N^{0.25})$. Since $(\alpha + 0.25) - (0.5\alpha + 1) = 0.5\alpha - 0.75 < 0$ as $\alpha < 1.5$, (14) follows.

We use $\mathcal{V}_{G_i}$ to denote the vertex set of G_i. Since $|G_{i,2}| \le \sum_{j\in\mathcal{V}_{G_i}} |G_j|$, we have $\sum_{i=1}^N |G_{i,2}| = \sum_{i=1}^N \sum_{j\in\mathcal{V}_{G_i}} |G_j| = \sum_{(i,j)\in G} (|G_i| + |G_j|) \le \sum_{e\in G} 2|A_e|$. So (14) ensures (17).

Also, $\sum_{e\in G} |A_e| = \sum_{(i,j)\in G} (|G_i| + |G_j| - 1) = \sum_{i=1}^N \sum_{j\in\mathcal{V}_{G_i}} |G_j| - |G| = \sum_{i=1}^N |G_i|^2 - |G|$ because for each i^*, $|G_{i^*}|$ appears $|G_{i^*}|$ times in the summation $\sum_{i=1}^N \sum_{j\in\mathcal{V}_{G_i}} |G_j|$. Since $|G| = o(N^{0.5\alpha+1})$ when $|G| = O(N^\alpha),\ \alpha < 1.5$, (14) and (16) are equivalent.

For (15), we have $\sum_{i=1}^N |G_i||G_{i,2}| \le \sum_{i=1}^N |G_i| \sum_{j\in\mathcal{V}_{G_i}} |G_j| = \sum_{i=1}^N \sum_{j\in\mathcal{V}_{G_i}} |G_i||G_j| =$

$2\sum_{(i,j)\in G}|G_i||G_j| \le 2\sum_{e\in A_e}|A_e|^2$. So (12) ensures (15).

Hence, all (11)-(18) can be derived from conditions in Theorem 5.1.

Next we prove result (ii). We have

$$\lim_{N\to\infty}\frac{\sigma_P^2}{\sigma_B^2} = \lim_{N\to\infty}\frac{|G| - \frac{mnN-2m^2-2n^2+2mn}{N^2(m-1)(n-1)}\left(\sum_{i=1}^N|G_i|^2 - \frac{4|G|^2}{N}\right) - \frac{2}{N(N-1)}|G|^2}{|G|}$$
$$= 1 - \lim_{N\to\infty}\left(\frac{\sum_{i=1}^N|G_i|^2}{N|G|} - \frac{4|G|}{N^2} + \frac{2|G|}{N^2}\right).$$

Because $|G| = O(N^\alpha),\ \alpha < 1.5$ and $\sum_{i=1}^N|G_i|^2 = o(N^{0.5\alpha+1})$ from the proof for result (i), we have $\lim_{N\to\infty}\left(\frac{\sum_{i=1}^N|G_i|^2}{N|G|} - \frac{4|G|}{N^2} + \frac{2|G|}{N^2}\right) = 0$, so $\lim_{N\to\infty}\frac{\sigma_P}{\sigma_B} = 1$.

Since $\mu_B - \mu_P = \frac{mn}{N^2(N-1)}|G|$, we have

$$\lim_{N\to\infty}\frac{\mu_B-\mu_P}{\sigma_P} = \lim_{N\to\infty}\frac{\sqrt{|G|}}{N-2} = 0, \tag{19}$$

when $|G| = O(N^\alpha),\ \alpha < 1.5$.

B Proof to Theorem 5.5

Let the density functions of the two multivariate distributions be f and g. When the similarity graph is a k-MST, $k = O(1)$, constructed on the Euclidean distance, following the approach in Henze and Penrose [1999], we have

$$\frac{R_1}{N} \to k\int\frac{\lambda^2 f^2(x)}{\lambda f(x)+(1-\lambda)g(x)}dx \quad \text{almost surely, and}$$
$$\frac{R_2}{N} \to k\int\frac{(1-\lambda)^2 g^2(x)}{\lambda f(x)+(1-\lambda)g(x)}dx \quad \text{almost surely.}$$

So

$$\frac{R_w - \mathbf{E}(R_w)}{N} = \frac{n}{N}\frac{R_1}{N} + \frac{m}{N}\frac{R_2}{N} - \frac{m(m-1)n/N + n(n-1)m/N}{N^2(N-1)}|G|$$
$$\to k\lambda(1-\lambda)\left(\int\frac{\lambda f^2(x)+(1-\lambda)g^2(x)}{\lambda f(x)+(1-\lambda)g(x)}dx - 1\right) \quad \text{almost surely.}$$

Since

$$\int\frac{\lambda f^2(x)+(1-\lambda)g^2(x)}{\lambda f(x)+(1-\lambda)g(x)}dx - 1 = \int\frac{\lambda f(x)(f(x)-g(x))}{\lambda f(x)+(1-\lambda)g(x)}dx$$
$$= \int\frac{(1-\lambda)g(x)(g(x)-f(x))}{\lambda f(x)+(1-\lambda)g(x)}dx,$$

we have

$$\begin{aligned}
&\int \frac{\lambda f^2(x) + (1-\lambda) g^2(x)}{\lambda f(x) + (1-\lambda) g(x)} dx - 1 \\
&= (1-\lambda) \int \frac{\lambda f(x)(f(x) - g(x))}{\lambda f(x) + (1-\lambda) g(x)} dx + \lambda \int \frac{(1-\lambda) g(x)(g(x) - f(x))}{\lambda f(x) + (1-\lambda) g(x)} dx \\
&= \lambda(1-\lambda) \int \frac{(f(x) - g(x))^2}{\lambda f(x) + (1-\lambda) g(x)} dx.
\end{aligned}$$

Therefore, $\int \frac{\lambda f^2(x) + (1-\lambda) g^2(x)}{\lambda f(x) + (1-\lambda) g(x)} dx \geq 1$ and it is strictly greater than 1 if f and g differ on a set of positive measure. Since $\sqrt{\mathbf{Var}(R_w)} = O(\sqrt{N})$ according to Theorem 3.1, the test is N^γ-consistent for any $\gamma > 0.5$.